

The Distributed Monitoring Framework

Brian Tierney, LBNL

Summary

The goal of the Distributed Monitoring Framework is to provide the ability to do performance analysis and fault detection in a Grid computing environment. This monitoring framework will provide accurate, detailed, and adaptive monitoring of all distributed computing components, including the network. Analysis tools will be able to use this monitoring data for real-time analysis, anomaly identification, and fault recovery. Services for network-aware applications will use this monitoring infrastructure in order to provide past, present, and future predictions of the network conditions, which the applications can then use to adapt their behavior. Applications such as widely distributed workflow management systems can then use the DMF as a generalized real-time event service.

The next generation of high-speed networks will allow DOE scientists unprecedented levels of collaboration. The focus of the DMF is to unify existing monitoring tools, provide seamless integration with a Grid environment, and to use these tools with scientific Grid applications. Experience gained will help determine what types of monitoring data best characterize various Grid applications, and how to reliably and efficiently disseminate that information to the proper subset of Grid middleware and end-user tools.

The DMF consists four main components: instrumentation, event publication, sensors, and event archiving. The term “event” is used here to mean a single timestamped unit of monitoring data. The current state of each component is described briefly below.

Instrumentation

Precision, real-time instrumentation of Grid applications and middleware is essential to understanding high performance data intensive applications. We have extended our previous work on the NetLogger Toolkit to provide the basis for this component. The NetLogger Toolkit included language APIs,

data collection tools, and a visualization tool called NLV.

To this we added a new, very efficient, binary logging format. We also added a fault tolerance mechanism, so that NetLogger is robust in the face of network instability. Finally, we added a "trigger interface" to NetLogger, that allows NetLogger to be turned on, off, or the logging level changed, in a running application.

Event Publication

The DMF bases its event publication on the architecture known as the Grid Monitoring Architecture (GMA), which was developed in coordination with other researchers in the Global Grid Forum (GGF). The GMA defines event “Producers”, event “Consumers”, and a distributed metadata “Registry” that allows Producers and Consumers to find each other. The event data goes directly from Producer to Consumer instead of through a central bottleneck.

We have written a prototype implementation of the GMA in the high-level Python language, called pyGMA, that uses standard Web Services protocols (the same ones used

in the newest release of the Globus Toolkit) for communicating control information between components, but which can use NetLogger as an efficient “data channel” for the monitoring event data.

We have combined the pyGMA and NetLogger “trigger” mechanism to allow remote activation of NetLogger instrumentation in running applications. We call this the NetLogger Activation Service, and have demonstrated it with an instrumented version of the High-Energy Physics framework Athena (part of the Atlas software group) at SC2002. This combination is exciting because it provides a relatively easy way to “peek into” a long-running process such as a physics calculation or large data transfer, without disrupting it.

Sensors

For sensors, the focus has been on integrating existing monitoring and sensors into the framework. Most of this work so far has been in standardization efforts.

Brian Tierney continued co-leading the "Glue Schema Work Group", which is tasked to define common schemas for interoperability between the EU physics grid projects (focusing on EDG and DataTag) and the US physics Grid projects (focusing in on PPDG, GriPhyN and iVDGL). Dan Gunter co-chaired the GGF Discovery and Monitoring Event Descriptions (DAMED) working group, which is defining a common level for describing “attributes” of, e.g., the GLUE schema. Brian also co-chaired the Network Monitoring Working Group, which is applying some of the principles of the DAMED work to the difficult problem of a precise vocabulary for network measurements.

Event Archiving

Historical analysis tools are necessary to understand why, for example, application performance “this week” is so much worse than it was “last week”. Condensing the data may mask important details, so this task requires an archive that can handle streams of detailed monitoring events. We designed and implemented a prototype of such an event archive using an open source relational database back-end (MySQL), and extensive buffering on disk to ensure that the database is never “frozen” by a large chunk of incoming events. The archive supports the GMA Consumer and Producer interfaces and has a web-based interface.

Next Steps

We are continuing our collaboration with the Atlas group, in particular with the Athena software group and Craig Tull, to add more sophisticated filtering and built-in analysis of the application monitoring to the Activation Service. Our goal is to diagnose problems in Athena, the Grid middleware, or the underlying systems in real-time and from anywhere in the world.

We are also extending pyGMA to include a truly distributed prototype of the GMA Registry, to allow easier location of monitoring results from, e.g., jobs scheduled on remote clusters. This also brings us closer to having a prototype general-purpose event service, which is needed by many Grid components. As always, we will coordinate and disseminate our progress in these areas through the GGF.

For further information on this subject contact:

Brian Tierney, Staff Scientist
Data Intensive Distributed Computing, DSD
Phone: 510-486-7381
blltierney@lbl.gov